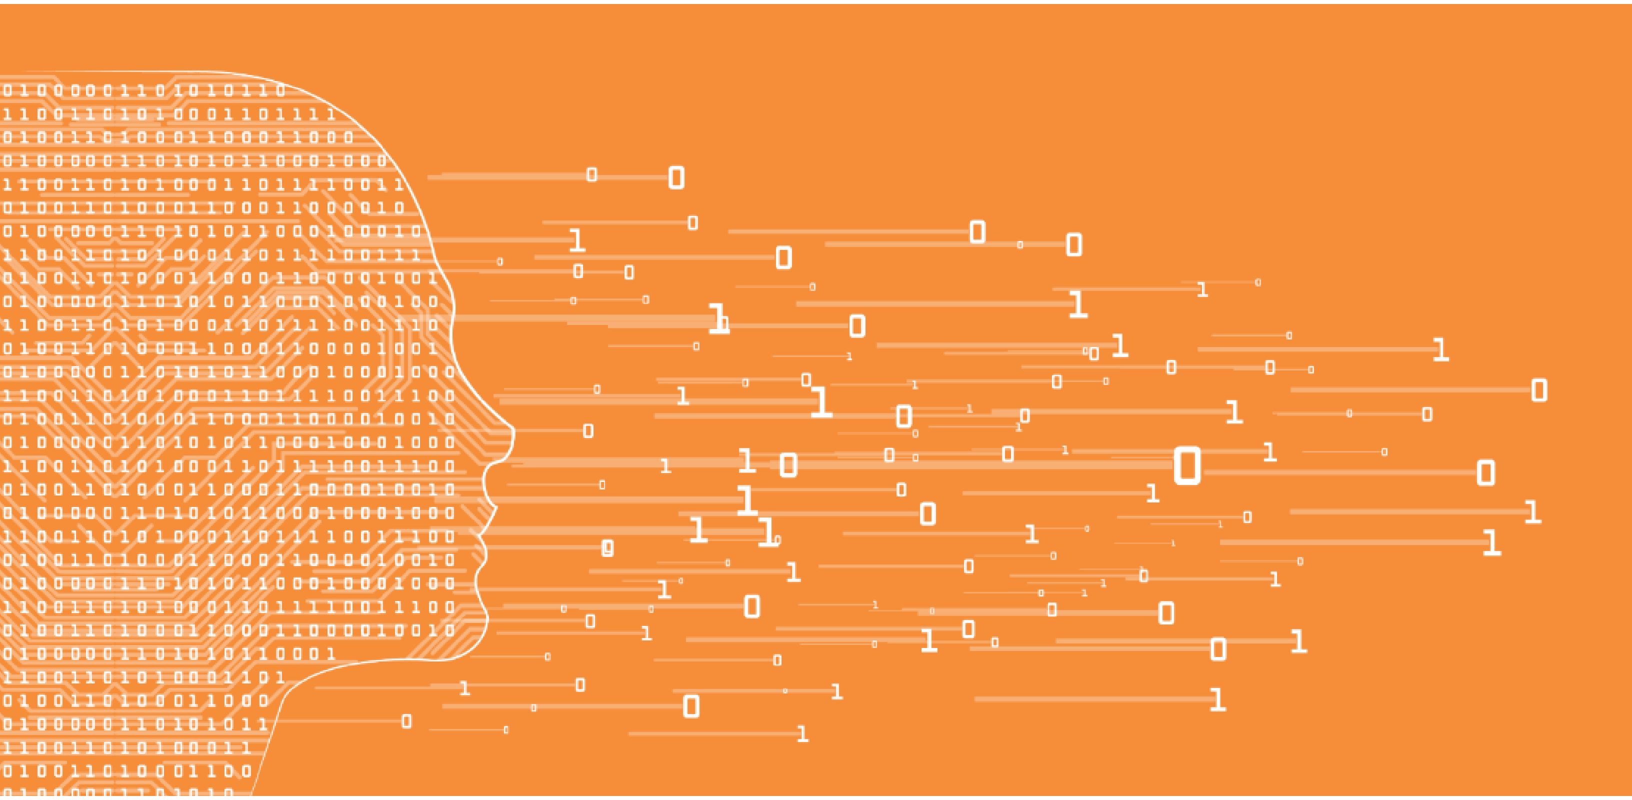

Clinithink

Not all healthcare AI is created equal: Accelerating insights and improving outcomes using clinical language models

Summary

The renewed popularity of AI makes many healthcare executives at once curious and dubious about the ways AI can help them work faster, gain precision, improve patient outcomes, and launch innovative new pursuits. Life science companies and healthcare providers want to know the best strategies for placing their AI investments. In this paper, we take a closer look at clinical natural language processing (CNLP), an alternative to solely using large language models (LLMs) when training AI to glean insights from healthcare data sources. When CNLP is paired with a powerful healthcare-specific ontology of terms like SNOMED CT, the results can yield more rapid, more targeted, and ultimately more valuable analysis. AI operations become more focused and transparent, making them easier to trace and refine, with granular insights that enable a faster, stronger patient cohort process while also improving clinical efficiency and documentation.

Contents

AI fervor yields AI fatigue. No one doubts the power and potential of modern AI, but many in healthcare are restless to learn more about how we can make it relevant to our business. Given our noble and often visionary goals, this concern is especially pointed: As AI advances, how do we best apply resources to make our institutions stronger and help people live healthier, longer lives?

As with most trends, a modicum of nuance helps dispel a mountain of generalities. This paper seeks to educate about the subtle but important differences between broad AI concepts in the mainstream versus a few targeted concepts that specifically benefit life sciences companies and healthcare systems. Primary among these is clinical natural language processing (CNLP), which makes free text computable within the healthcare domain. It also introduces opportunities for precision not generally afforded by large language models (LLMs) that simply scan huge datasets looking for previously trained phrases. At Clinithink, for example, our implementation of CNLP, called CLiX, defies the general-to-specific decision pathways of other language models and instead zeroes in on terms and phrases contained in the deeply modeled SNOMED CT ontology, a comprehensive and multilingual schema of clinical healthcare terms that organizations use globally to exchange clinical information.

SNOMED CT is a rich, detailed global standard for healthcare data. Where LLMs can be a blunt instrument for parsing and repurposing clinical language, CNLP using SNOMED CT offers a much-needed scalpel. By focusing on the highly specialized nature of clinical language, CNLP benefits from using SNOMED CT as a framework, accelerating the generation of high-quality clinical insights. This speed and accuracy, combined with other key advantages including increased granularity and transparency, enables AI for healthcare data scientists that is ultimately far more useful and effective than other AI methods.

These concepts are especially illuminating for researchers, data scientists, and healthcare IT decision makers who seek to maximize their AI investments but need help breaking down the key differentiators between AI types and methods. Here, we'll cover a bit of the background on CNLP and how it's situated within the spectrum of modern AI techniques. Then, we'll talk through the areas where CNLP provides unique—indeed, unprecedented—access to vital insights that can help reshape clinical, research, and financial operations across our industry, with a clear target of improving lives and expanding the scope of treatment options.

How does SNOMED CT optimize CNLP?

To maximize the efficacy and usefulness of CNLP, SNOMED CT organizes concepts into an *acyclic taxonomic* (IS-A) hierarchy. Imagine SNOMED CT as a large oak tree, with a root, a trunk, and various branches, eventually ending with leaves at the tip of each branch. It has direction (upward) while also being acyclic (i.e. it does not loop back on itself), and it is hierarchical in that the branches have split points or nodes. Each node leads to other nodes, and the nodes eventually terminate at the leaves of the tree.

By moving outward in a single direction, this acyclic taxonomic structure creates an "inheritance hierarchy" where each node (split point) inherits all the attributes of the nodes above it, with each predecessor node (parent, grandparent, great-grandparent, and so on) passing down its traits to the child nodes that follow it.

In this example, "lung disease" expresses the general condition, while "pneumonia," "infectious pneumonia," and "viral pneumonia" express more specific instances of a lung disease, and all relevant attributes of the grandparent (lung disease) term are passed down to the child and grandchild terms.

But there's one important exception to our tree example. Whereas a leaf on a tree only ever connects to a single branch, acyclical taxonomic hierarchies can have multiple parents above them, allowing nodes to inherit more than one set of attributes—for example, "Infectious pneumonia" is also a child node of "Infectious disease." This useful multiple-inheritance principle is the key to understanding why SNOMED CT makes an ideal ontology for CNLP: Because each leaf can have more than one parent, each term can inherit a plentiful complexity of associations without losing its place in the hierarchy.

The key strength of this taxonomic structure is the richness it brings to the overall ontology, allowing data to be recorded and later accessed at many different levels of aggregation. SNOMED CT concepts are linked by approximately 1,360,000 links, called relationships, and these allow SNOMED to express over 1.5 billion concepts through all combination of elements.

For example, SNOMED CT associates the four terms related to lung disease in the following hierarchy

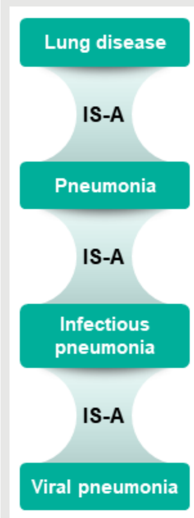

The beauty—and necessity—of a "narrow/deep" AI focus

The current AI hype cycle is awash with controversy and cautionary tales about bots, deep fakes, and the automation of knowledge work. These are all compelling, even alarming, topics. Lost in this discussion, however, is the simple fact that within all this newfound computing power lives the potential to dramatically change society for the better. In the case of CNLP, this amounts to better population health, clinical research, and business outcomes through the strategic application of AI in the healthcare and life sciences markets.

If you've read about AI, you likely already know about LLMs, those gargantuan, trainable neural networks that compute human-like artificial language and images based on vast sets of input and statistical training. LLMs are currently popular for their ability to generate content, giving us generative AI, a relatively new capability within the 70 years or so that AI has been around.

CNLP also yields language, but in a very different way from most LLMs. With a narrow focus on the language set within a single domain—healthcare—CNLP analyzes, develops, and refines language for use in the distinct knowledge area of clinical settings. It starts with an ontology like SNOMED CT as its definitive terminology schema and refines its recognition capability by being trained on thousands upon thousands of known permutations and associations of short healthcare terms and phrases. This intense spotlight enables a faster path to accuracy and discovery for healthcare insights, especially as we continue training it and enriching our AI capabilities for healthcare as we go along.

To see the difference, consider this example. When treating bacterial infections, doctors don't prescribe broad-spectrum antibiotics for every case; they choose targeted antibiotics based on the specific bacteria involved. This approach not only treats the infection more effectively but also reduces the risk of resistance. Similarly, while popular LLMs start with a broad range of language associations, CNLP proceeds from a specialized subset of clinical terms, leading to more precise and relevant insights. Just as a doctor uses the right antibiotic for a specific infection, CNLP uses targeted language models to deliver more accurate analysis.

Focused language + clinical context = next-level healthcare insights

So, CNLP succeeds by performing its AI in an industry-focused area, and derives its underlying knowledge of terminology from powerful, specialized resources like SNOMED CT. But how do we get from there to genuinely new levels of healthcare insight? The answer is complicated, but it starts with context. The AI you use needs to account for multiple clinicians saying related things but using different terms to do so. Finding correlation and relevance requires recognizing how these dissimilar terms are related. By aggregating and analyzing healthcare data on this deeper level, you can more accurately trust an AI system to find new pathways to relevant clinical intelligence.

Achieving this deeper analysis is not an easy task. Success comes from looking beyond electronic medical record (EMR) data and incorporating the valuable information found in free-text documents, such as providers' dictated or hand-typed notes. These unstructured data elements, unlike digital data, are not easily analyzed by technology and are often misunderstood by LLM-based AI systems.

One significant reason for this is that most LLMs lack the semantic information model used by CNLP. This model excels at associating a wide variety of terms and phrases within a specific, narrow context. As a result, CNLP provides healthcare organizations across the provider–pharma spectrum with insights from unstructured data and the opportunity for driving unprecedented medical advancements. It overcomes the limitations of LLM-based analysis.

A key benefit of CNLP is its ability to handle negation, uncertainty, temporal obfuscation (such as past versus present state), and other complexities in the text as seen in phrases like:

- "the patient denied breathlessness"
- "there was no evidence of edema"
- "possible history of dysphagia"

By referencing a comprehensive ontological schema like SNOMED CT, CNLP can readily recognize and associate phrases by using domain-specific knowledge, greatly enhancing accuracy and relevance. In contrast, the generalized methods of LLM analysis are much more likely to miss these important associations.

To observe the importance of an AI understanding phrases on this level, consider the following example:

> Mother has a severe allergy to penicillin, but she has no known penicillin allergy.

In this scenario, many LLMs are challenged to distinguish between the mother, who has a severe allergy to penicillin, and the patient, who does not. If they fail this challenge, meaning is negated or lost. By design, CNLP is better suited to untangle these nuances and associations.

Maximizing business and clinical benefits in healthcare AI

Using an AI system that rapidly gleans high-quality, highly relevant intelligence can expand your business benefits and clinical reach immensely. For example, Clinithink participated in a recent American Society of Clinical Oncology (ASCO) study where we used CLiX to enable proactive identification of early lung cancer. At this level, with this approach, the impossible becomes possible when the technology supporting you is trained to understand clinical nuances that it can then parse at unprecedented speed.

And the practical uses go on from there. CNLP, as an exemplar of focused, hyper-contextual AI, can:

- Enable **cohort recruitment specialists** to more quickly identify, assemble, and review highly specific patient cohorts for clinical research and post-market activities
- Help **clinicians and researchers** identify types of cancer at earlier stages than ever before, at scale

- Reduce service denials and other **revenue cycle management** pitfalls through automation
- Empower **health systems** to tackle healthcare inequities while enhancing clinical productivity

Let's take a brief look at each of these opportunities in more detail.

BETTER PHARMA COHORTS, FASTER

Building successful pharma cohorts requires identifying and qualifying the subset of the patient population most likely to benefit from a new medicine or label. This process is typically challenging, expensive, and time-intensive, costing valuable research time and money that could otherwise be spent directly on other discovery and development activities that contribute to population health and revenue goals.

The essential data needed to accelerate and improve accuracy for cohort assembly lies deep in clinical records. With CNLP, an AI system can apply critical intelligence across vast amounts of unstructured data buried within patients' medical histories. The speed and precision of CNLP over other AI methods enable it to correlate the right information for making new cohorts more strategic and impactful. And it can do this in a matter of hours, versus the months or even years that pharma companies and the cohort recruitment specialists they work with can spend trying to attain the same results using other techniques. In one example from our experience at Clinithink, CNLP was used to identify patients who met characteristics of lung disease, possibly related to lung cancer. CNLP made it possible to rapidly identify a cohort of thousands, from a population of millions, saving valuable clinical researcher time that otherwise would have been spent looking at records manually and with brute force tools.

CNLP, when correctly applied, can also benefit life sciences companies in other ways, including pinpointing the critical signals that drive research requirements, more accurately anonymizing data to protect patient privacy, and accelerating early detection of patients at risk for a specific condition or disease.

BOOSTING CLINICAL PRODUCTIVITY

Health systems face ever-growing demand for services, which they continuously must reconcile with rising operational costs and workforce capacity risks like turnover and burnout. Better data analysis enables smarter use of a system's resources, both human and financial, by helping providers and administrators perform more efficiently and find newer, healthier ways of achieving positive patient outcomes.

By applying powerful, healthcare-focused technologies like CNLP, health systems can streamline operations in ways that help providers deliver more sustainable and targeted interventions. They also gain a greater advantage to be proactive by using insights from unstructured data like providers' notes to take preventive steps with patients, which reduces healthcare costs in chronic disease populations while increasing the level of timely care and attention these populations need.

CNLP even makes it possible to identify populations that demonstrate common disease conditions, to better treat or even prevent the onset of the condition. At Barts Hospital, part of the NHS in the UK, providers used CNLP to better identify patients with a risk of diabetic foot ulcer. CLiX scanned 14.2 million documents to find patients with diabetic foot disease, identifying 30% more patients with diabetes and 375% more patients with diabetic foot problems compared to the manual review process. These findings provide immense benefit to patients as well as significant cost savings to the hospital system.

QUICK STATS

14.2 million documents to find patients with diabetic foot disease

Identifying **30%** more patients with diabetes

375% more patients with diabetic foot problems

IMPROVING CLINICAL DOCUMENTATION FOR PROVIDERS AND ISVS

Data is also revolutionizing payment processing in healthcare, including not just providers but also the independent software vendors (ISVs) who support them with technology and services for revenue cycle management (RCM). The better these companies can manage and automate patient data, the more exacting and impactful they can be in RCM activities.

The vectors on which providers scale up these efficiency gains all pertain to clinical documentation and how it's handled. By readily identifying gaps in documentation and applying business rules based on these gaps to the underlying automation logic, CNLP offers a way to shorten revenue cycles and ease the burden on teams who manage it.

As part of clinical documentation improvement, CNLP also offers opportunities for provider maturity in its denials and appeals workflows. The acute analysis CNLP performs on documents, including the valuable insights it can find buried in unstructured data like clinical notes, empowers providers to efficiently process denials and generate appeal letters in a fraction of the time otherwise required for manual review—and with higher levels of accuracy—by identifying criteria in patient accounts that make them more likely to be denied for common conditions. At Northwell Health System in New York, CNLP helped the clinical documentation improvement (CDI) short-stay team gain significant efficiency overall, saving valuable staff time while increasing documentation and billing accuracy.

Evaluating clinical language models and addressing challenges

When opting for an AI that uses LLM or CNLP, it's important to understand key differences in the ways each method addresses various success factors pertinent to healthcare data analysis. The Clinithink blog series on this topic dives into details about many of these differences. For now, let's look at a summary of what you should consider.

TRACKING PROVENANCE WITHIN AI WORKFLOWS

One vital asset of using natural language processing in healthcare is the transparency it affords its users. With CNLP, context is king, and transparency is its scientific realm. Data scientists, like any other scientists, place a high value on knowing the source of the intellectual raw material they're using to conduct experiments that will hopefully yield meaningful results. For healthcare, this means being able to look at the outcome from an AI prompt and trace back its logical precepts and calculations, so that the data scientists can track down the sources of unexpected or erroneous results and retrain the language model as needed.

For most LLMs, maintaining this kind of provenance on any level is all but impossible—or, if it is possible, it's impractical due to the additional processing costs required. LLMs often lack the real-world understanding to grasp semantics as they process information, relying instead on pattern recognition and word associations. Contrastingly, CNLP uses robust methodologies designed to inherently understand and capture semantics, history, negation, family references, and countless other nuances within the healthcare content it processes. By training on this extensive repository and incorporating human expertise to review and refine their model, data scientists can base AI projects on a superior understanding of typical clinical grammar patterns.

ACHIEVING GRANULARITY, OVERCOMING BIAS, AND REDUCING FALSE POSITIVES

Providers often use slightly different words when referring to the same thing. This fact greatly complicates the goal of deriving insights from patient histories, because it requires intelligently equivocating what one specialist says so it can be correlated to another's notes. The more granular an AI system can be in this work, the more we can trust the output to be consistent and reliable. Implementing a clinical language model in any type of healthcare process requires the extraction of extremely precise detail so that synonymous terms and concepts can be correlated successfully.

Many LLMs are far too slow to meet these granularity requirements and many other healthcare intelligence needs. When you try to speed up the LLMs, processing costs increase once again and, worse, the LLMs will usually lose critical granularity along the way, making them inadequate for obtaining high-value insights. Superior clinical natural language processing (CNLP) can greatly enhance accuracy and relevance of their output by recognizing and associating phrases like these by referencing domain-specific knowledge tied to an ontology like SNOMED CT.

Another limitation of LLMs is their known ability to inadvertently fabricate details that are not present in the source data or prompt. CNLP helps to exclude these hallucinations, prioritizing fact and truth over guesswork or faulty interpolation. And bias during AI model training can warp the AI model's output, negatively affecting efficacy, outcomes, and medical decision making throughout a patient's treatment journey.

Choosing a CNLP provider

The traditional computer science approach to natural language processing, including CNLP, has been centered on mathematically-based “brute force” operations on LLMs using machine learning (ML). At Clinithink, we feel clinical text is too complicated to be processed successfully using this method. As such, our approach does not principally rely on machine learning but rather part-of-speech (POS) tagging, lemmatization, and similar linguistic constructs to isolate noun-phrases in the input text. Here’s how our patented approach works:

1. CLiX breaks up the input text into chunks that carry the POS tags, and within milliseconds, a pipeline of techniques involving more than 30 different parallel tasks operates on each chunk.
2. The input is then mapped to a specialized proprietary language model, resulting in a highly codified and computable output format—a fully post-coordinated SNOMED CT expression— that preserves the granular detail represented by the input narrative and conveys the value.
3. These expressions are then abstracted, leveraging the structure and description logic of SNOMED CT to resolve what was found in the source text into a potential match, or not, for the concept—for example, a phenotype or detailed clinical characteristic—as is needed to create business value in the relevant use case.

The content abstracted by CLiX ultimately matches the details doctors capture very frequently in a medical record that provide essential additional context, and that make perfect sense to other humans but completely confound most AI systems. Our experience indicates that ML-based approaches (including LLMs) will struggle to manage the complexity of this domain.

The true business opportunities in healthcare AI are shaped by selecting an AI partner whose approach most benefits your healthcare goals and your bottom line. Using CNLP with SNOMED CT provides a robust foundation for extracting valuable, granular insights from unstructured data in highly usable, transparent ways. The future revolutions in healthcare lie within the context of doctors' notes and other non-EMR sources. With the right method, you can access these insights swiftly, accurately, and with the specificity needed to drive meaningful advancements. By harnessing unstructured data effectively, you unlock new business opportunities and innovations that propel healthcare forward.

Clinithink LLC

www.clinithink.com

E info@clinithink.com